

AI Disclosure Statement — HoneyMoney

Prepared by: Team HoneyMoney · Last updated: 10 July 2026

HoneyMoney is an AI-native financial wellness platform for Malaysian households. We believe transparency about *how* and *where* we use artificial intelligence is a condition of trust — especially in a product that handles people's money. This statement discloses, honestly and completely, the AI in our product, the AI we used to build it, and the limitations and safeguards we have put in place. Where a claim is an estimate rather than a measured fact, we say so.

1. AI Inside the Product

What AI we use. HoneyMoney uses **swappable AI providers** — Google **Gemini Flash**, **Groq**, and local **Ollama** — selected per deployment, with a per-call token and cost ledger for governance. A multimodal model performs two jobs:

- **OCR / data extraction.** When a user forwards an e-wallet or receipt screenshot (Touch 'n Go eWallet, Maybank MAE, GrabPay, ShopeePay, and paper receipts) to our Telegram bot, the model reads the image and extracts structured fields: merchant, amount, date, and a suggested category.
- **Insight generation.** Our AI persona, "Honey", turns a user's stored transaction history into plain-language, proactive financial guidance.

What data the model processes. For extraction, the model receives the forwarded image and returns structured text. For insights, it receives *derived, minimized* financial data (aggregated figures and category summaries) needed to answer the user's question — not a raw dump of a user's entire ledger. With the local Ollama option, this processing can stay entirely on-device. We do not send images or data to the model for advertising, model training, or any purpose beyond serving the user's request.

Human-in-the-loop confirmation. HoneyMoney never silently records what the AI reads. Every parsed transaction is shown back to the user for **explicit confirmation** before it is written to their ledger. Users can correct the merchant, amount, category, or bucket. The AI proposes; the human decides.

The 3-Bucket model is rules-based, not AI-decided. Allocation of money across Fixed Non-Negotiables, the Future Shield auto-savings bucket, and private personal wallets follows user-configured rules. AI helps users *understand* their money; it does not autonomously move or lock funds.

2. Privacy and PDPA Handling

- **Raw images are discarded.** After extraction and user confirmation, the original forwarded screenshot is not retained. We keep the structured transaction record, not the picture of the receipt.
- **Data minimization.** We store the minimum needed to run the service — nodes, edges, and transactions in our **local-first PocketBase knowledge graph** — and send the model only what a given task requires.
- **Local-first storage.** Data lives on the user's/team's own machine rather than a central cloud database, which is both a privacy safeguard and a genuine PDPA story.
- **Marital-safe by design.** Private personal wallets are private. Honey is explicitly constrained never to expose one partner's private-wallet detail to the other.
- **PDPA alignment.** We operate consistent with Malaysia's Personal Data Protection Act 2010: users consent to processing, data is used only for the stated purpose, and users can request correction or

deletion of their records.

3. AI Used to Build HoneyMoney

We used AI coding assistants (Anthropic's **Claude** and **Claude Code**) during development for scaffolding, boilerplate, refactoring, and drafting code across our stack. This accelerated delivery, but **every AI-assisted contribution was reviewed, tested, and accepted by a human engineer** before it entered the product (stack: Next.js 16 PWA, local-first PocketBase, Telegram Bot API, and multi-provider AI). Architecture decisions, security-sensitive logic, and financial rules were designed and verified by our team. AI was a tool in the workshop, not the author of record.

4. Limitations and Safeguards

We want judges and users to understand what our AI *cannot* guarantee:

- **OCR is not 100% accurate.** Model extraction can misread amounts, merchants, or dates — particularly on low-quality images or unfamiliar receipt formats. This is precisely why mandatory human confirmation exists: no unverified figure enters a user's ledger.
- **No automated, binding financial advice.** Honey offers general guidance and behavioural nudges for financial wellness. It does **not** provide licensed financial, investment, tax, or legal advice, and it does not execute transactions on a user's behalf. Guidance is framed as suggestion, not instruction.
- **AI can be wrong or incomplete.** Generated insights may occasionally be generic or mistaken. Users remain in control of all decisions about their money.
- **Data minimization as a standing safeguard.** Limiting what we collect and what we send to the model reduces both privacy risk and the blast radius of any error.

5. Our Commitment

If any statement here becomes inaccurate as the product evolves, we will update this disclosure. We would rather understate our capabilities and be believed than overstate them and mislead. Candor is not just a competition requirement for us — it is the only way to earn the trust required to help families manage their money.